IN THE SPECIFICATION

Please replace the following paragraphs with the amended paragraphs below.

Paragraph beginning on page 5, lines 17-23:

The disadvantages discussed above are overcome and additional features are disclosed by a method for streaming content striped in RAID 5 format from an array of disk drives to a plurality of subscribers to minimize disruptive service from a disk drive failure. The method includes accessing content data on an extent-by-extent basis from a plurality of disk drives in an array and ~~streams~~streaming the content data to the plurality of subscribers on an extent-by-extent basis, sequentially, from the plurality of disk drives.

Paragraph beginning on page 7, lines 20-28:

The disadvantages of the RAID 3 and RAID 5 data striping formats are overcome, and synergistic benefits are achieved by implementing various modes of operation by using a "stream regeneration" algorithm, a "data regeneration" algorithm, and a "recovery carousel-serving" algorithm, which together combine the advantageous features of both RAID 3 and RAID 5 disk drive accessing formats when a disk drive in a RAID 5 array of disks fails. Once a stream server initiates these ~~two~~ algorithms, disk drive failures are masked from the perspective of subscribers receiving streamed content by continuing to serve content in a reduced stream capacity mode.

Paragraph beginning on page 8, line 23 to page 9, line 9:

The video server 310 contains a Statistical Disk Scheduler (SDS) 370, a central processing unit (CPU) 314, and memory element 317. The SDS 370 is coupled to the plurality of disk drives (hereinafter "disk drives") 320 in each disk array 319 by paths 330_0 through 330_k (collectively paths 330) (e.g., fiber-channel), and to the memory 317 by data path 377. The video server 310 sends access requests to a particular disk drive array 319 (loop server) through the video switch 304 and along paths 330 to disk drives 320, where each disk drive 320 has its own internal queue 325_1 through 325_k

(collectively queues 325) for buffering the access requests. Data read from the disk drives 320 is transmitted back to the video server 110 along paths 330. The paths 330 are serially coupled (e.g., "daisy chained") to form the loop 331 (e.g., a fiber-channel loop). Preferably, the system 300 employs multiple loops ~~131~~331 to interconnect the disk drives 320 in the disk drive arrays 319 (loop server), such that the data transfer rate amongst the disk drives 320 and the video server 310 is increased over that of the single loop system 331. For a detailed understanding of an exemplary method for transferring data between independent fiber channel loops in an interactive information distribution system, the reader is directed to commonly assigned U.S. patent ~~application Serial No. 09/458,320, filed December 10, 1999~~6,766,393, which is incorporated herein by reference in its entirety.

Paragraph beginning on page 9, lines 10-25:

In one embodiment, an *SDS Selection Procedure* is utilized to select requests from the three SDS queues (not shown) and forward the requests to an associated disk drive queue 325_n located within each of the disk drives 320_n. The three SDS queues may include a Steady-State queue (SSQ), a non-steady-state queue (NSQ), and an optional other request queue (ORQ). The SDS Selection Procedure uses worst-case access times, request priorities, and time deadlines in determining which request to forward to the disk drive queue. The general strategy of the SDS Selection Procedure is to select a NSQ request only when such a selection will not cause any of the SSQ requests to miss their time deadlines, even if the NSQ request and all requests in the SSQ were to take their worst-case access times. If such a guarantee cannot be made, then the first request in the SSQ is always selected. For a detailed understanding of an exemplary statistical disk scheduler and the SDS queues, the reader is directed to commonly assigned U.S. patent ~~application Serial No. 09/268,512, filed March 12, 1999~~6,378,036, which is incorporated herein by reference in its entirety.

Paragraph beginning on page 12, lines 17-33:

Where a disk drive 320 has failed, the server 310 transitions from the carousel-serving algorithm 334 to the stream regeneration algorithm 339. The basic concept underlying the stream regeneration algorithm 339 is to switch from the RAID 5 format of reading each extent 402 one at a time, to the RAID 3-like format of reading the entire parity group simultaneously. In particular, the bad (i.e., failed) disk drive is no longer accessed, however, all the extents across the other disk drives 320 are read in the parity group 404. The parity information 410 is then used to reconstruct (regenerate) the missing true data from the defective disk drive. From this point on, realtime parity correction can be performed to correct failed reads from the defective disk drive. Furthermore, rebuilding of the data on the disk drive array 319 can proceed as well. Specifically, the data regeneration algorithm 336 executes a disk drive rebuild, and in one embodiment the lost data is written in the spare extents 410 on the disk drives 320 in a disk drive array 319 as a low priority task, as discussed below. Once all the data is written to the spare extents, the server 310 can transition to the recovery-carousel-serving algorithm 338 to stream the contents to the subscribers 360.

Paragraph beginning on page 13, lines 18-27:

In the embodiment shown in FIG. 4, spare extents 410 are distributed across the disk drives 320 in the disk drive array 319. ~~Alternately~~Alternatively, the RAID 5 striping may be provided across the disk drives 320 of the array 319 without dedicating spare extents 410 in a parity group 404. In any of the embodiments (i.e., predictive or non-predictive disk failures, with/without spare extents), the stream regeneration algorithm, data regeneration algorithm, and recovery carousel-serving algorithm may be utilized, either singularly or in combination, to stream content to subscribers and rebuild lost data in the event a disk drive 320 fails, as is discussed in detail below with regard to the embodiments of FIGS. 5 through 7.

Paragraph beginning on page 14, lines 11-17:

At step 502 (drawn in phantom), a determination is made as to whether a disk drive 320 has failed. If, at step 502, the determination by the server 310 is answered negatively, then the method 500 proceeds to step 504, where the server 310 continues to operate in the normal mode of operation using the carousel-serving algorithm 334 and ~~continuing~~continues to monitor for disk drive failures. If, however, at step 502, the determination by the server 310 is answered positively, then the method 500 optionally proceeds to step 506 (drawn in phantom).

Paragraph beginning on page 14, line 18 to page 15, line 4:

~~In an~~At optional step 506, content loads onto ~~the disk drive array 319 having the~~a failed disk drive 320 of a disk drive array 319 are normally disallowed. Although a disk drive 320 may have failed in a disk drive array, in some instances content (e.g., movies, trailers, and/or any other type of content) may still be loaded onto the operable disk drives in the failed disk drive array 319. In particular, the server 310 calculates the parity data for the data to be loaded ~~in~~from the parity extents of the drive. As such, the parity data for the additionally loaded data is available and may be used to regenerate data onto a replacement disk drive at a later time. Such instances where additional~~Additional~~ content loads may be performed ~~in the case where~~occur when there are only a few disk drive arrays 319 coupled to a server 310, there is a high demand for some particular content (e.g., a new movie), and replacement of the failed disk drive may not occur within an acceptable time period. However, in most instances, ~~the~~ high demand content is distributed (e.g., load balanced) amongst numerous disk drive arrays 319 (i.e., loop servers) to accommodate subscribers requesting particular high demand content. Furthermore, disallowing content loads during the failed disk drive mode of operation reduces the overall bandwidth requirements, thereby increasing available bandwidth for higher priority tasks, such as streaming content to subscribers or disk rebuilds. As such, optional step 506 is typically implemented to disallow additional content loads onto the failed disk drive array 319.

Paragraph beginning on page 15, lines 5-14:

At optional step 508 (drawn in phantom), some of the subscribers accessing the disk drives 320 in the failed disk drive array 319 are migrated to (~~i.e.,~~and served by) a non-failed disk drive array 319. By migrating some of the subscribers receiving content from the failed disk drive array 319 to another disk drive array, the number of streams generated by the failed disk drive array is reduced. The subscribers remaining on the failed disk drive array 319 are then able to receive increased bandwidth and higher priority for streaming content ~~and~~to help mask the disk drive failure. Additionally, reducing the stream demand on the failed disk drive array increases the bandwidth of the disk drives 320 during the data reconstruction phase of the method 500, as discussed below.